HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P. O. Box 272400
Fort Collins, Colorado 80527-2400

PATENT APPLICATION

ATTORNEY DOCKET NO. 10001666-1

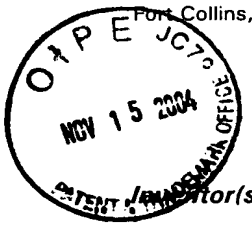

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Blaine D. Gaither

Confirmation No.:

Application No.: 09/703,428

Examiner: Marc D. Thompson

Filing Date: 10/31/2000

Group Art Unit: 2144

Title: Fault tolerant storage system having an interconnection fabric that also carries network traffic

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Sir:

Transmitted herewith in ***triplicate*** is the Appeal Brief in this application with respect to the Notice of Appeal filed on ***09/10/2004***.

The fee for filing this Appeal Brief is (37 CFR 1.17(c)) $330.00.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

( ) (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d) for the total number of months checked below:

( )	one month	$110.00
( )	two months	$420.00
( )	three months	$950.00
( )	four months	$1480.00

( ) The extension fee has already been filled in this application.

(X) (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account ***08-2025*** the sum of ***$330.00***. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

(X) I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, Alexandria, VA 22313-1450. Date of Deposit: ***Nov. 10, 2004***

OR

( ) I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number_____ on _____

Number of pages:

Typed Name: ***Joanne Bourguignon***

Signature: _____

Respectfully submitted,

Blaine D. Gaither

By _____

Robert W. Bergstrom

Attorney/Agent for Applicant(s)
Reg. No. ***39,906***

Date: ***Nov. 10, 2004***

Telephone No.: ***206.621.1933***

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent application of:

Inventors: Blaine D. Gaither
Serial No.: 09/703,428
Filed: October 31, 2000
For: Fault Tolerant Storage System having an Interconnection Fabric that also Carries Network Traffic

Examiner: Marc D. Thompson
Group Art Unit: 2144
Docket No. 10001666-1
Date: November 10, 2004

APPEAL BRIEF

Commissioner of Patents and Trademarks
Washington, DC 20231

Sir:

This appeal is from the decision of the Examiner, in an Advisory Action mailed on August 10, 2004, finally rejecting claims 1-36.

REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

RELATED APPEALS AND INTERFERENCES

Applicants' representative has not identified, and does not know of, any other appeals of interferences which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-36 are pending in the application. Claims 1-36 were finally rejected in the Office Action dated April 21, 2004. Applicants' appeal the final rejection of claims 1-36, which are copied in the attached CLAIMS APPENDIX.

STATUS OF AMENDMENTS

No Amendment After Final is enclosed with this brief. The last Response was filed June 21, 2004.

SUMMARY OF CLAIMED SUBJECT MATTER

Overview

Applicants' system uses a fault-tolerant-storage system ("FTSS") to store and forward network communications packets between two nodes, or host computers. In general, network communications systems involve network communications media that, directly or indirectly, transmit messages, or packets, often through intermediary routing nodes, from a source computer to a destination computer. Communications media include such well-known communications media as the Ethernet and other local area networks, fibre channel, various types of wide area networks ("WANs"), and other types of communications media. Applicants have recognized that there may be significant available data-transfer bandwidth between computers and FTSSs that, although until now employed only for data transfer operations, such as file transfer and database management system operations, can be alternatively modified for transfer of network packets from a source computer to a destination computer, both the source computer and the destination computer connected to the FTSS through dedicated fibre-channel links, or other types of dedicated communications media, as an alternative to network communications pathways linking the source and destination computers. This method involves storing packets received by the FTSS from source

computers in highly available storage within the FTSS, and subsequently retrieving the stored packets and forwarding the retrieved packets by the FTSS to destination computers. Normally, storing and forwarding packets by an FTSS in such a manner would be considered far too inefficient and incurring far greater latency than normally acceptable for network communications, but, as recognized by Applicants, when network communications is bandwidth limited, or fails, the store-and-forward method through an FTSS provides alternative bandwidth or an available, alternative pathway for exchanging of packets between source and destination computers, or nodes. Moreover, because of the speed of the dedicated links, and high-throughput design of certain FTSS systems, the transmission efficiency may, in fact, be acceptable.

Independent Claim 1

Claim 1 relates to transmitting network packets from source nodes to destinations through a fault-tolerant storage system by transmitting a packet over an FTSS interconnection fabric, or dedicated FTSS communications link, from the source node to the FTSS, storing the packet by the FTSS in highly reliable fault-tolerant storage media, and subsequently transmitting the packet from the FTSS to the destination node.

Dependent Claims 2 – 9

Dependent claims 2-9 provide elaboration of elements of claim 1, including transmission of acknowledgement messages by the FTSS, packet management by the FTSS, packet protocol encapsulation, and other details.

Independent Claim 10

Independent claim 10, like claim 1, relates to transmitting network packets from source nodes to destinations through a fault-tolerant storage system by transmitting a packet over an FTSS interconnection fabric, or dedicated FTSS communications link, from the source node to the FTSS, storing the packet by the FTSS in highly reliable fault-tolerant storage media, and subsequently transmitting the packet from the FTSS to the destination node.

Dependent Claims 11 – 16

Dependent claims 11-16 provide elaboration of elements of claim 10, including transmission of acknowledgement messages by the FTSS, identification of various types of fault-tolerant storage, packet protocol encapsulation, and other details.

Independent Claim 17

Independent claim 17 relates to transmitting network packets from source nodes to destinations by transmitting the packet from the source nodes through an external network to an FTSS, storing the packet by the FTSS in highly reliable fault-tolerant storage media, and subsequently transmitting the packet from the FTSS to the destination node by an FTSS interconnection fabric.

Dependent Claims 18 – 24

Dependent claims 18-24 provide elaboration of elements of claim 10, including transmission of acknowledgement messages by the FTSS, identification of various types of fault-tolerant storage, packet protocol encapsulation, and other details.

Independent Claim 25

Independent claim 25 relates to an interconnection system that transmits network packets from source nodes to destinations by transmitting the packet from the source nodes through an external network to an FTSS, storing the packet by the FTSS in highly reliable fault-tolerant storage media, and subsequently transmitting the packet from the FTSS to the destination node by an FTSS interconnection fabric.

Dependent Claims 26 – 34

Dependent claims 11-16 provide elaboration of elements of claim 10, including transmission of acknowledgement messages by the FTSS, identification of various types of fault-tolerant storage, packet protocol encapsulation, and other details.

Independent Claim 35

Independent claim 35 relates to a server in an interconnection system that transmits network packets from source nodes to destinations by transmitting the packet from the source nodes through an FTSS interconnection fabric using a file I/O interface.

Independent Claim 36

Independent claim 25 relates to an FTSS that represents the hub of an interconnection system that transmits network packets from source nodes to destinations by transmitting the packet from the source nodes through an external network to an FTSS, storing the packet by the FTSS in highly reliable fault-tolerant storage media, and subsequently transmitting the packet from the FTSS to the destination node by an FTSS interconnection fabric.

GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. In an Office Action dated April 21, 2004 ("Office Action"), the Examiner rejected claims 1-36 under 35 U.S.C. § 103(a) as being unpatentable over Utter et al., U.S. patent No. 5,815,649 ("Utter") in view of Byers et al., U.S. Patent No. 5,809,543 ("Byers").

ARGUMENT

ISSUE 1

1. Whether claims 1-36 are unpatentable over Utter in view of Byers.

Utter discloses a fault-tolerant computer system in which user terminals intercommunicate through a network. In fact, networked communication in which users interact with various terminals interconnected with one another and with one or more large computer systems is well known. Such systems have been commercially available for at least 30 years. However, as Applicants' representative's has pointed out in a previously filed Response, Utter makes no mention or suggestion that information transferred between two user terminals is intercepted by the fault-tolerant computer system and stored in fault-tolerant storage media.

The Examiner references lines 55-62 of column 4 of Utter in support of the

unjustified conclusion that the network shown in Utter transfers data from one user terminal to another using a fault-tolerant-computer-system-based store-and-forward technique. The entire paragraph is provided, below:

> In addition to transferring requests from the user terminal 12(d) to the fault-tolerant computer system 11, the networks 13(n) may also download data from the fault-tolerant computer system 11 to the user terminals 12(d) for local processing by the user terminals 12(d) or for display to a user on a display terminal which may be provided with respective ones of the user terminal 12(d). In addition, the networks 13(n) may transfer data from the user terminals 12(d) to the fault-tolerant computer system 11 for processing or storage, and in addition may transfer data among the user terminals 12(d). In one embodiment, the networks 13(A) and 13(B) are each in the form of conventional high-speed (11 Mb/second) Ethernet networks, which transfer information in the form of messages. As is conventional in Ethernet networks, messages generated by one device connected to a network 13(A), 13(B) (that is, by a user terminal 12(d) or by the fault-tolerant computer system 11) contain information to be transferred, as well as an address which identifies the intended recipient or recipients of the information. (emphasis added)

It is manifestly clear that a conventional network, such as an Ethernet network, is being discussed in this paragraph. Utter clearly and correctly discusses the fact that the conventional network can be used for request/response interactions between a user interacting with the fault-tolerant computer system. This is well known, and has been well known for at least 25 years. Utter correctly observes that the conventional network may also transfer data among the user terminals. This is also well known. Utter even explains that all of such network transactions are based on the fact that one computer, whether a terminal or the fault-tolerant computer system, can address a message to any other computer on the network. If a source terminal wishes to send a message to a destination terminal on a network, the source terminal composes the message and transmits the message to the destination terminal by appropriately addressing the message. Terminals and computers on the network other than the destination terminal do not receive the message, unless specially configured to do so for network analyzer applications. Utter does not teach or suggest that Utter's fault-tolerant computer system is used as a network analyzer. Ethernet controllers within terminals and computer systems receive all packets sent through an Ethernet, and filter the received packets so that the terminals and computer systems that include the Ethernet controllers receive only those messages that are addressed to them. These aspects of Ethernet communications systems are well known, and described, in detail, in numerous textbooks and websites.

The Examiner appears to state that Utter teaches or suggests, in order to implement user-terminal-to-user-terminal communications, that Utter's fault-tolerant

computer system abandons well-known Ethernet convention and well-known principles of computer networking, and rather than simply allowing the user terminals to directly transmit messages to one other, Utter's fault-tolerant computer system instead forces user terminals to transmit messages to a storage controller (31(a) or 31(s)) within the fault-tolerant computer system, which stores the message in fault-tolerant storage, later retrieves the message from fault-tolerant storage, and then forwards the retrieved message to the recipient user terminal. The Examiner's proposed system would slow messaging in Utter's system by several orders of magnitude, overburden the Ethernet controllers associated with storage nodes, and make the fault-tolerant computer system an unnecessary bottleneck for messaging. There is no suggestion in Utter that Utter contemplated such a system, and no teaching, mention, or suggestion that user-terminal-to-user-terminal communication involves anything other than a simple transmission of messages from one user directly to another user via the Ethernet. Instead, Utter teaches a *conventional Ethernet network*, as explicitly stated by Utter, in which a user terminal can directly transmit a message to another user terminal without intervention by the fault-tolerant computer system. Indeed, in the above paragraph, Utter indicates that the network may transfer data among the user terminals as well as transfer data from user terminals to the fault-tolerant computer system.

There is also no indication in Utter that network packets are stored in fault-tolerant storage for any reason. Instead, as explicitly stated by Utter beginning on line 57 of column 4:

> Similarly, the storage nodes 30(s) perform the data processing services for the user terminals 12(d) as described above in connection with the fault-tolerant computer system 11, *and the networks 13(n), switches 18(h) and networks 20(k) serve to transfer storage and retrieval requests from the user terminals 12(d) and processing nodes 16(m) to the storage nodes 30(s) which are to execute the request and return any data and status information that is to be returned from the storage nodes 30(s) executing the request.* (emphasis added)

Thus, Utter makes it clear that network activity is devoted to transferring data storage and retrieval requests from a processing node to a storage node and from the storage node to the processing node. Not once does Utter mention or suggest a processing node sending a network message to a storage node, which stores the network message, and then forwards that network message to another processing node.

In Applicants' clearly claimed method and system, a source computer transmits a message to a destination computer by: (1) transmitting a network packet to an

FTSS; (2) storing the network packet in a fault-tolerant storage medium by the FTSS; and (3) forwarding the network packet from the FTSS to the destination computer. This is not a conventional networking approach. In a conventional network, such as that described by Utter, the source computer directly transmits the network packet to the destination computer. Conventional network operation is highly desirable from fault-tolerance and computational efficiency standpoints. An Ethernet is a bus exactly in order to avoid making any particular node a single point of failure and a bottleneck for communications. The Examiner seems to suggest Utter employs a much older, star-topology, store-and-forward network, in which a central computer stores and forwards all messages transmitted between points of the star. *But, even in the old star-topology systems, fault-tolerant storage would not have been used by the central computer, because fault-tolerant storage generally requires that multiple copies of data be stored, which, in a networking system, would generally be considered to be unnecessary and extremely inefficient.* Almost all modern communications systems, even ring-based systems, avoid store-and-forward messaging, because they are considered to be, in general, slow, computationally inefficient, and extremely brittle from a reliability standpoint. Nowhere in Utter does Utter suggest a store-and-forward communications system. Instead, Utter discloses a conventional networking system in which a source node transmits a message directly to a destination node.

Byers discloses an "outboard file cache extended processing complex for use with a host data processing system for providing closely coupled file caching capability" (Abstract). The cited passages of Byers describe various hardware components of the internal buses and links within Byers file cache and between the file cache and host computers, but neither disclose nor suggest a first host computer sending a network protocol message to a second host computer through the outboard file cache. Indeed, all the host computers are directly interconnected via communications links, as clearly shown in the Figure shown on the first page of Byers. Byers is simply unrelated to Applicants' claimed invention.

In an Advisory Action dated August 8, 2004 ("Advisory"), the Examiner responded to Applicants' arguments. First, the Examiner states:

> Examiner disagrees with the Applicant assertion that the claims necessarily require a store-and-forward type arrangement as argued. The breadth of the claims recite the transfer and subsequent storage of a packet received from a source node within a fault tolerant storage system, where the packet is stored, and then transferred to a destination node, thereby effecting a transfer of the packet(s) from the source node to

the destination node(s) through an intermediate, reliable, fault-tolerant storage subsystem.

Claim 1 is provided below, with emphasis added:

> 1. A method of transmitting a network packet from a source node to a destination node, wherein the source and destination nodes are coupled to a fault tolerant storage system (FTSS) via an FTSS interconnection fabric, the method comprising:
>
> transmitting the packet from the source node to the FTSS via the FTSS interconnection fabric;
>
> storing the packet in highly reliable fault-tolerant storage media of the FTSS; and
>
> transmitting the packet from the FTSS to the destination node via the FTSS interconnection fabric.

In the Current Application, a networked system in which the claimed invention is practiced is shown in Figure 4, and described in the paragraph beginning on line 11 of page 9. In particular, Figure 4 clearly shows server computers 80, 82, 84, 86, 88, and 90 interconnected to an FTSS 92 via an FTSS interconnection fabric 94, and an external network 96 that interconnects the FTSS and servers by traditional networking methods. As stated in the Current Application, "the terms 'node' and 'server' may be used interchangeably" (Current Application, page 9, lines 17-18). Claim 1 clearly states that a packet is transmitted from the source node to the FTSS, stored in highly reliable fault tolerant storage, and transmitted by the FTSS to the destination node. Claim 1 includes a first step carried out by a source node, a second step carried out by the FTSS, and a third step carried out by the FTSS, resulting in transmission of a packet to a destination node. The source and destination nodes are clearly defined and described in the Current Application, and are not within the FTSS. Furthermore, the FTSS receives a packet, stores the packet in fault-tolerant storage, and transmits, or forwards, the packet to a destination node. This is, without doubt, a clearly described and clearly claimed store-and-forward-based method for packet transmission. The FTSS both stores and forwards the packet. Utter's system does not operate in this manner. In Utter's system, packets are transmitted directly from source terminals to destination terminals via a conventional Ethernet, and there is no suggestion in Utter that network packets are intercepted, stored, and forwarded by the FTSS. Such a method would require specialized initialization of Ethernet controllers by the FTSS and specialized logic in the FTSS to receive, store, and forward packets not addressed to the FTSS. In fact, because the Ethernet is a bus, the FTSS would have no way of preventing a message sent from a source terminal to

a destination terminal from being directly received by the destination terminal. At best, the FTSS could intercept the message, and then resend it, resulting in the destination terminal receiving the message twice, and probably becoming thoroughly confused. There is no suggestion in Utter that any of these specialized initializations and logic are implemented or used. Applicants' representative respectfully asserts that there is no basis for the Examiner's statement either in the Current Application, the current claims, or Utter.

Next, the Examiner states in the Advisory Action:

> The teachings of Utter provide, minimally, the storage and retrieval of data for processing at various types of processing end points. See, inter alia, Column 3, Lines 17-67. The provision for storage of information (one or more packets) from source device(s), stored within the storage subsystem, and the reception of this data from the storage subsystem at a destination device(s), was disclosed. That is, the use of a terminal to store information in the fault tolerant storage system, the storage, and then subsequent retrieval of that information from a(nother) terminal in the system, meets the claimed invention. Thus, the invention as claimed in the currently overly broad independent claims, as presented, are met by the combination of teachings of Utter and Byers. Applicant seems to argue that the art fails to teach basic store-and-forward type transfer of information. Minimally, since the processing units disclosed by Utter take input information for storage from the terminal(s), store the information, then retrieve the information at other/same terminal(s), this simplistic overview does not make sense.

Applicants fully appreciate the fact that information can be sent by a terminal to the FTSS for storing the information in the FTSS. Applicants also fully appreciate that the same terminals, or other terminals, can separately retrieve stored information from the FTSS. That is how FTSSs operate. These are two-way transactions between one terminal and the FTSS. Basic FTSS operation is not at all the subject mater to which claim 1 is directed. The referenced portion of Utter's disclosure discusses various types of user terminals, and the operation of a conventional Ethernet communications medium that interconnects the terminals with the FTSS. The referenced passage simply discusses the fact that terminals can request storage and retrieval of data of the FTSS, and that the data is exchanged via the Ethernet in a conventional manner. The data in the referenced portion of Utter is not one or more network packets. Instead, the data is encoded within one or more network packets. Network packets include headers that contain, among other things, Ethernet addresses of intended recipients, as stated even in the referenced portion of Utter on lines 62-67:

> As is conventional in Ethernet networks, messages generated by one device connected to a network 13(A), 13(B) (that is, by a user terminal 12(d) or by the fault-tolerant computer system 11) contain information to be transferred, as well as an address which identifies the intended recipient or recipients of the information.

The messages discussed in this passage are the packets discussed in the Current Application and in the current claims. Utter clearly describes transfer of information in Ethernet packets to and from the FTSS, but does not discuss storing and forwarding of the *Ethernet packets themselves*. Storing and retrieving information from an FTSS by a client computer or terminal in 2-way transactions issued through a communications system has nothing whatsoever to do with a method for storing and forwarding network-level packets by an FTSS that are intended to be sent by a first computer external to the FTSS and received by a second, different computer external to the FTSS. Utter clearly states that Utter's system uses a convention Ethernet, in which packets are directly addressed to their intended recipients, and not a specially implemented system in which packets sent by a first terminal to a second terminal are instead intercepted by the FTSS, stored by the FTSS, and separately forwarded to the intended recipient. In fact – such a store-and-forward method is nearly impossible on the Ethernet, since an Ethernet controller in any devices interconnected to the Ethernet receives at least all packet headers of all packets, in order to decide whether or not the packet is addressed to the Ethernet controller, and, if addressed to the Ethernet controller, receives and stores the packet into memory. Store-and-forward techniques, by contrast, can occur when two links are interconnected by a routing node, or in serial media in which packets are passed along a chain of nodes. The Ethernet is a bus – not a node-to-node link.

Finally, the Examiner's statement is replete with reference to "overly broad claims." For example, the Examiner states:

> "Thus, the invention as claimed in the currently overly broad independent claims ..."
>
> "It is noted that Applicant has failed to modify the claim language to distinguish over the prior art of record by clarifying and substantially narrowing the claim language."
>
> "Failure for Applicant to significantly narrow definition/scope of the claims and supply arguments commensurate in scope with the claims implies the Applicant intend broad interpretation be given to the claims."

Applicants' representative has noticed a trend in Office Actions received from the USPTO for Examiners to refuse to carefully consider arguments and the cited art unless applicants first amends claims. Applicants are entitled by law to claim as broadly as the knowledge in the art allows. Applicants are in no way obligated to amend claims unless the Examiner provides relevant references on which Applicants' claims read or that makes the claimed subject matter obvious, presuming, of course, that the Applicants have originally crafted claims to cover that

which they believe to be novel. The Examiner, in the current case, has failed to do provide such references. Utter's system discloses terminals and an FTSS interconnected by a conventional FTSS, and nothing of any relevance to Applicants' claimed invention. Byers is completely unrelated. Applicants clearly disclose and claim a system in which an FTSS stores and forwards packets between external source and destination computers. Neither Utter, Byers, nor Utter and Byers in combination teach, suggest, or mention a store-and-forward communications method in which an FTSS stores and forwards network packets. As a result, the 35 U.S.C. § 103(a) rejections of claims 1-36 are both legally and technically unjustified. The Examiner, in the present case, cannot decide that claims are overly broad *per se*, but must, instead, identify prior art on which the claims read, or that makes the claimed subject matter obvious.

CONCLUSION

Applicants clearly disclose and claim a system in which an FTSS stores and forwards packets between external source and destination computers. Utter's system discloses terminals and an FTSS interconnected by a conventional FTSS, and nothing more of relevance to Applicants' claimed invention. Byers is completely unrelated to the claimed invention. Neither Utter, Byers, nor Utter and Byers in combination teach, suggest, or mention a store-and-forward communications method in which an FTSS stores and forwards network packets. Claims 1-36 are therefore not made obvious under 35 U.S.C. § 103(a) by Utter, Byers, or Utter and Byers in combination.

Applicants respectfully submit that all statutory requirements are met and that the present application is allowable over all the references of record. Therefore, Applicants respectfully request that the present application be passed to issue.

Respectfully submitted,
Blaine D. Gaither
OLYMPIC PATENT WORKS PLLC

By _____
Robert W. Bergstrom
Reg. No. 39,906

Olympic Patent Works PLLC
P.O. Box 4277
Seattle, WA 98104
206.621.1933 telephone
206.621.5302 fax

CLAIMS APPENDIX

1. A method of transmitting a network packet from a source node to a destination node, wherein the source and destination nodes are coupled to a fault tolerant storage system (FTSS) via an FTSS interconnection fabric, the method comprising:

transmitting the packet from the source node to the FTSS via the FTSS interconnection fabric;

storing the packet in highly reliable fault-tolerant storage media of the FTSS; and

transmitting the packet from the FTSS to the destination node via the FTSS interconnection fabric.

2. The method of claim 1 and further comprising:

sending an acknowledgment from the FTS S to the source node guaranteeing delivery of the packet to the destination node.

3. The method of claim 1 and further comprising:

sending an acknowledgment from the destination node to the FTSS acknowledging receipt of the packet at the destination node.

4. The method of claim 1 and further comprising:

deleting the packet from the highly reliable fault-tolerant storage media after the packet has been transmitted to the destination node.

5. The method of claim 1 and further comprising:

retaining the packet in a packet queue of the FTS S for a period of time after the packet has been transmitted to the destination node.

6. The method of claim 1 wherein the packet is transmitted from the source node to the FTSS, and from the FTSS to the destination node as a file I/O transaction

7. The method of claim I wherein the packet is transmitted from the source node to the FTSS, and from the FTSS to the destination node by encapsulating the packet in a protocol of an interface used to implement the FTSS interconnection fabric.

8. The method of claim I wherein storing the packet in highly reliable fault-tolerant storage media of the FTSS comprises storing the packet in a nonvolatile write cache of the FTSS.

9. The method of claim 1 wherein storing the packet in highly reliable fault-tolerant storage media of the FTSS comprises storing the packet in a redundant array of independent disks of the FTSS.

10. A method of transmitting a network packet from a source node to a destination node, wherein the source node is coupled to a fault tolerant storage system (FTSS) via an FTSS interconnection fabric, and the destination node is coupled to the FTSS via an external network, the method comprising:

transmitting the packet from the source node to the FTSS via the FTSS interconnection fabric;

storing the packet in highly reliable fault-tolerant storage media of the FTSS; and transmitting the packet from the FTSS to the destination node via the external network.

11. The method of claim 10 and further comprising:

sending an acknowledgment from the FTSS to the source node acknowledging that the FTSS has received the packet, but delivery of the packet to the destination node will be attempted but can not be guaranteed.

12. The method of claim 10 and further comprising:

deleting the packet from the highly reliable fault-tolerant storage media after the packet has been transmitted to the destination node.

13. The method of claim 10 wherein the packet is transmitted from the source node to the FTSS as a file I/O transaction.

14. The method of claim 10 wherein the packet is transmitted from the source node to the FTSS by encapsulating the packet in a protocol of an interface used to implement the FTSS interconnection fabric.

15. The method of claim 10 wherein storing the packet in highly reliable fault-tolerant storage media of the FTSS comprises storing the packet in a nonvolatile write cache of the FTSS.

16. The method of claim 10 wherein storing the packet in highly reliable fault-tolerant storage media of the FTSS comprises storing the packet in a redundant array of independent disks of the FTSS.

17. A method of transmitting a network packet from a source node to a destination node, wherein the destination nodes is coupled to a fault tolerant storage system (FTSS) via an FTSS interconnection fabric, and the source node is coupled to the FTSS via an external network, the method comprising:

transmitting the packet from the source node to the FTSS via the external network;

storing the packet in highly reliable fault-tolerant storage media of the FTSS; and transmitting the packet from the FTSS to the destination node via the FTSS interconnection fabric.

18. The method of claim 17 and further comprising:

sending an acknowledgment from the destination node to the FTSS acknowledging receipt of the packet at the destination node.

19. The method of claim 17 and further comprising:

deleting the packet from the highly reliable fault-tolerant storage media after the packet has been transmitted to the destination node.

20. The method of claim 17 and further comprising:
retaining the packet in a packet queue of the FTSS for a period of time after the packet has been transmitted to the destination node.

21. The method of claim 17 wherein the packet is transmitted from the FTSS to the destination node as a file I/O transaction.

22. The method of claim 17 wherein the packet is transmitted from the FTSS to the destination node by encapsulating the packet in a protocol of an interface used to implement the FTSS interconnection fabric.

23. The method of claim 17 wherein storing the packet in highly reliable fault-tolerant storage media of the FTSS comprises storing the packet in a nonvolatile write cache of the FTSS.

24. The method of claim 17 wherein storing the packet in highly reliable fault-tolerant storage media of the FTSS comprises storing the packet in a redundant array of independent disks of the FTSS.

25. A networked system comprising:
a plurality of nodes;
a fault tolerant storage system (FTSS); and
an FTSS interconnection fabric coupling the plurality of nodes to the FTSS; wherein each node includes:
a network protocol stack for processing network I/O;
an interface for sending data to and receiving data from the FTSS interconnection fabric; and
a packet conversion unit for linking the network protocol stack, thereby allowing network traffic to flow between the node and the FTSS via the FTSS interconnection fabric;
and wherein the FTSS includes:
nonvolatile fault-tolerant storage media for storing data;

a file operations unit for completing file 110 operations to the nonvolatile fault tolerant storage media; and

a network routing agent for receiving packets from source nodes of the plurality of nodes, storing packets in the nonvolatile fault-tolerant storage media, and transmitting packets to destination nodes of the plurality of nodes.

26. The networked system of claim 25 wherein the network routing agent of the FTSS sends an acknowledgment to the packet conversion unit of the source node guaranteeing delivery of the packet to the destination node.

27. The networked system of claim 25 wherein the packet conversion unit of the destination node sends an acknowledgment to the networking routing agent of the FTSS acknowledging receipt of the packet at the destination node.

28. The networked system of claim 25 wherein the network routing agent of the FTSS deletes the packet from the nonvolatile fault-tolerant storage media of the FTS S after the packet has been transmitted to the destination node.

29. The networked system of claim 25 wherein the network routing agent of the FTSS retains the packet in a packet queue stored on the nonvolatile fault-tolerant storage media for a period of time after the packet has been transmitted to the destination node.

30. The networked system of claim 25 wherein packets are transmitted from source nodes to the FTSS, and from the FTSS to the destination nodes as file I/O transactions.

31. The networked system of claim 25 wherein packets are transmitted from source nodes to the FTSS, and from the FTSS to the destination nodes by encapsulating packets in a protocol of an interface used to implement the FTSS interconnection fabric.

32. The networked system of claim 25 wherein the nonvolatile fault-tolerant storage media of the FTSS includes a nonvolatile write cache.

33. The networked system of claim 25 wherein the nonvolatile fault-tolerant storage media of the FTSS comprises a redundant array of independent disks.

34. The networked system of claim 25 and further comprising an external network coupled to the FTSS, wherein the FTSS routes packets between the plurality of nodes coupled to the FTSS, and external nodes coupled to the external network.

35. A server comprising:
a plurality of applications and system utilities;
an interface for coupling the server to a fault tolerant storage system (FTSS); and I/O system services coupled between the plurality of applications and system utilities and the interface, the 110 system services including:
a file system for processing file I/O operations between the plurality of applications and system utilities and an FTSS via the interface; and
a network protocol stack for processing network packets between the plurality of applications system and utilities and other network nodes; wherein the network protocol stack links into the file system to carry network packets to an FTSS via the interface.

36. A fault tolerant storage system (FTSS) comprising:
nonvolatile fault-tolerant storage media for storing data;
a file operations unit for completing file I/O operations to the nonvolatile fault tolerant storage media; and
a network routing agent for receiving packets from source nodes coupled to the FTSS via an FTSS interconnection fabric, storing

packets in the nonvolatile fault-tolerant storage media, and transmitting packets to destination nodes via the FTSS interconnection fabric.

EVIDENCE APPENDIX

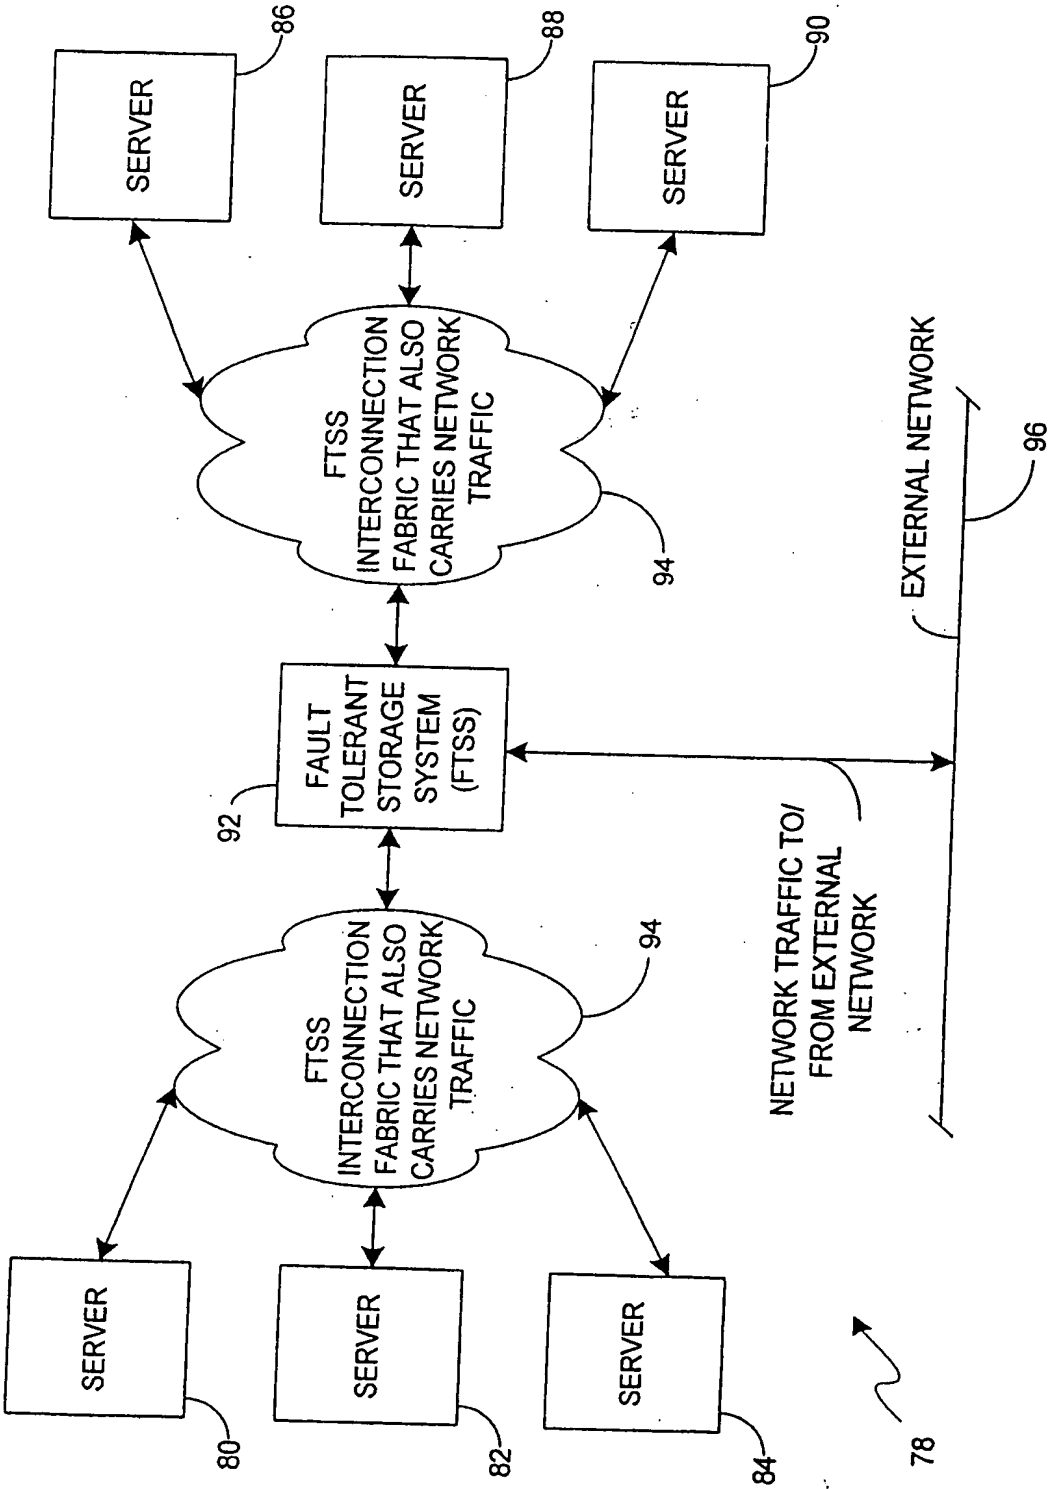

FIG. 4

RELATED PROCEEDINGS APPENDIX

None.